




WHITE PAPER

Liquid cooling

THE 2026 IMPERATIVE



EN

04_26

TABLE OF CONTENTS

1 INTRODUCTION	p. 4	9.1 CURRENT STATE: 400 V AC DISTRIBUTION	p. 21
2 THE EVOLUTION OF LOAD IN DATA CENTERS	p. 5	9.2 800 V DC DISTRIBUTION	p. 21
2.1 FROM THE "DENSITY ERA" TO THE "POWER ERA"	p. 5	9.3 LIQUID-COOLED PSU	p. 21
2.2 THE EVOLUTION OF TDP IN NVIDIA PROCESSORS	p. 5	9.4 CHALLENGES OF 800 V: SAFETY AND ENVIRONMENTAL MANAGEMENT	p. 21
2.3 THE PHYSICAL LIMIT OF AIR: THE "THERMAL WALL"	p. 5	10 CONCLUSION	p. 22
2.4 CONCLUSION	p. 5	11 APPENDIX A	p. 23
3 THERMODYNAMICS OF LIQUID COOLING	p. 6	11.1 CHEMICAL MANAGEMENT OF THE FLUID AND MATERIAL COMPATIBILITY	p. 23
3.1 AIR VS LIQUID	p. 6	12 APPENDIX B	p. 24
3.2 CONDUCTION AND CONVECTION	p. 6	12.1 OPERATIONAL COMMISSIONING AND TESTING PROTOCOL	p. 24
3.3 THE ROLE OF PG25 IN DIRECT-TO-CHIP	p. 6		
3.4 FLUIDS FOR TWO-PHASE DIRECT-TO-CHIP AND IMMERSION COOLING	p. 6		
4 COOLING TECHNOLOGIES	p. 7		
4.1 SINGLE-PHASE DIRECT-TO-CHIP (DTC)	p. 7		
4.2 TWO-PHASE DIRECT-TO-CHIP (DTC)	p. 8		
4.3 SINGLE-PHASE IMMERSION COOLING (IC)	p. 9		
4.4 TWO-PHASE IMMERSION COOLING (IC)	p. 10		
5 LIQUID COOLING ECOSYSTEM	p. 11		
5.1 HYBRID ROOM ARCHITECTURES	p. 11		
5.2 THE COMPLETE ECOSYSTEM OF LIQUID COOLING	p. 14		
5.3 LIQUID DISTRIBUTION WITHIN THE ROOM	p. 14		
6 COOLANT DISTRIBUTION UNIT	p. 16		
6.1 CDU COMPONENTS	p. 16		
6.2 REGULATION LOGIC	p. 17		
6.3 NVIDIA SERVER COMPLIANCE	p. 17		
7 EFFICIENCY, SUSTAINABILITY, AND HEAT RECOVERY	p. 18		
7.1 BEYOND PUE: THE DEFINITION OF ERF	p. 18		
7.2 HEAT RECOVERY STRATEGIES	p. 18		
7.3 REGULATORY FRAMEWORK	p. 18		
8 OPERATIONAL CHALLENGES, MAINTENANCE, AND SAFETY	p. 20		
8.1 CHEMICAL MANAGEMENT AND MONITORING OF PG25	p. 20		
8.2 ROUTINE MAINTENANCE OF THE CDU	p. 20		
8.3 SERVER MAINTENANCE	p. 20		
9 FUTURE FRONTIERS – 800 V DATA CENTERS	p. 21		

1 INTRODUCTION

The data center industry is undergoing the most significant transformation of the past thirty years. For decades, server rooms have been designed and built according to an air-first approach: raised floors, suspended ceilings, and aisle containment systems have enabled the optimization of air-based cooling and the full exploitation of its potential. However, the advent of Generative Artificial Intelligence and high-performance computing (HPC) has introduced thermal loads that exceed the physical limits of forced-air convection, pushing the data center industry to a true breaking point.

In 2026, with processors having far exceeded the threshold of 100 W of TDP (Thermal Design Power), the question is no longer “if” liquid cooling should be adopted, but rather “how” to implement it in a simple, scalable manner that is integrated with the existing infrastructure.

This White Paper aims to guide industry professionals through the transition from air cooling to liquid cooling, a shift that extends beyond the physical medium used for heat removal and impacts the entire room ecosystem: from manifolds and piping to quick couplings, and up to pressure and temperature control within the circuits.

KEY POINTS

- **Overcoming the “Thermal Wall”:** analysis of the reasons why traditional infrastructures based on CRAC/CRAH systems are no longer technically sustainable for densities exceeding 40 kW per rack;
- **Liquid cooling technologies:** technical in-depth analysis of direct-to-chip solutions and immersion systems, including evaluation of architectures, operating principles, and application scenarios;
- **Coolant Distribution Unit (CDU):** analysis of the key components, operating principle, and the regulation and control logics governing the distribution of the coolant fluid;
- **Measurable sustainability:** assessment of the impact of liquid cooling on PUE reduction, with values approaching 1.1, and on heat recovery opportunities for district heating applications.

2 THE EVOLUTION OF LOAD IN DATA CENTERS

2.1 From the “Density Era” to the “Power Era”

Historically, data center design has been driven by the optimization of airflow: over the past decade, 5 kW racks have gradually been replaced by increasingly high-performance solutions, reaching up to 15 kW per rack. This transition was made possible by maximizing the potential of air cooling, with the introduction of hot/cold aisle containment systems and the elevation of raised floors to reduce pressure losses and create a uniform static pressure plenum serving the room.

However, between 2019 and 2020, the advent of High Performance Computing (HPC) led to a gradual yet irreversible shift in the rules of the game: HPC servers increased density to the point of reaching 40 kW per rack. With loads of this magnitude, air ceases to be an efficient cooling solution, as the velocities required to remove heat from the servers would generate turbulence levels that render the process inefficient.

To date, we have exceeded 100 kW per rack and are preparing to reach 200 kW per rack: air cooling, in its traditional form, is no longer a viable option.

2.2 The evolution of TDP in NVIDIA processors

The primary driver of this acceleration is the evolution of Thermal Design Power (TDP).

Non più di 10 anni fa, una CPU di fascia alta operava nell'intervallo 150-200W. Oggi la situazione è cambiata radicalmente:

- **NVIDIA V100/A100:** 300/400 W per module. It marked the first true entry of 40 kW/rack into the commercial sector;
- **NVIDIA H100 (Hopper):** TDP up to 700 W per module;
- **NVIDIA B200 (Blackwell):** TDP reaching 1000–1200 W per single GPU.

When 8 of these GPUs are integrated into a single server node (HGX), the heat generated within just a few rack units is comparable to that of an entire data center row from ten years ago. Dissipating 10 kW from a single 4U server using fans would require airflow rates and rotational speeds high enough to generate vibrations harmful to components and sound power levels exceeding 100 dB(A).

Generation	Year	TDP Max	Recommended cooling system
P100 (Pascal)	2016	300W	Air cooling (traditional)
V100 (Volta)	2017	300W	Air cooling (traditional)
A100 (Ampere)	2020	400W - 500W	Air cooling (rear-door)
H100 (Hopper)	2023	700W	Direct-to-Chip Liquid Cooling
B200 (Blackwell)	2025	1000W - 1200W	Direct-to-Chip Liquid Cooling

2.3 The physical limit of air: the “Thermal Wall”

The concept of the “thermal wall” emerges when the energy cost and the space required to move air exceed the cooling benefit itself. Three factors make air cooling obsolete for AI workloads:

- **Volumetric heat capacity:** water has a volumetric heat capacity approximately 3,400 times higher than air; therefore, to dissipate even large amounts of heat from a server, the flow rate circulating through a pipe with a diameter of only a few millimeters is sufficient;
- **Approach:** to cool 1000 W chips with air, very low supply temperatures (18–20°C) would be required, minimizing the possibility of leveraging free cooling. Liquid cooling, on the other hand, enables operation with a reduced temperature approach and, consequently, higher supply temperatures;
- **Fan energy consumption:** the power absorbed by fans increases cubically with respect to airflow rate: in high-density air-cooled racks (40 kW per rack), fans can account for up to 20% of the server’s total power consumption.

2.4 Conclusion

From a design standpoint, exceeding 40 kW per rack renders air an obsolete medium: the transition to liquid cooling becomes mandatory.

3 THERMODYNAMICS OF LIQUID COOLING

3.1 Air vs Liquid

The superiority of liquid cooling over air cooling is not quantitative, but qualitative.

From a thermodynamic standpoint, the most significant parameter for heat transfer efficiency is the **Volumetric Heat Capacity [kJ/m³·K]**. That of water is more than **3,600 times higher** than that of air: therefore, for the same amount of heat transported, the required liquid flow rate will be much lower than that of air.

To understand the spatial impact of this concept, let us analyze what is required to remove **100 kW**, with a 10°C deltaT, from a single rack:

- **Air cooling:** a flow rate of approximately 30,000 m³/h would be required, which is clearly impractical to implement;
- **Liquid cooling (PG25):** approximately 9 m³/h are required, equivalent to 2.5 L/s. This flow rate can be managed with piping of only a few centimeters in diameter, drastically reducing the infrastructural footprint.

3.2 Conduction and convection

Traditional air cooling relies on forced convection, a process constrained by the extremely low thermal conductivity of air, which effectively acts as a natural insulator.

In liquid cooling, and specifically in direct-to-chip (DTC) systems, heat is removed through copper cold plates. This approach enables:

- **Minimize thermal resistance:** heat is transferred from the processor die to the fluid through a high-conductivity metallic path, minimizing thermal gradients;
- **Enable “warm water cooling”:** thanks to the efficiency of liquid, it is possible to operate with high supply temperatures (35°C–45°C) while still keeping processors within their operating limits. This makes it possible to leverage free cooling in regions where it was historically unfeasible, eliminating dependence on chillers and drastically reducing PUE (Power Usage Effectiveness).

3.3 The role of PG25 in direct-to-chip

Although pure water is an excellent heat transfer fluid, its use in a data center is not recommended due to the risks of corrosion, biological growth, and freezing. The industry standard for DTC has consolidated around PG25 (a mixture of 25% propylene glycol and 75% demineralized water with inhibitors).

PG25 represents the ideal engineering compromise:

1. **Corrosion protection:** chemical inhibitors protect secondary loops composed of different metals (copper, stainless steel), preventing galvanic corrosion;
2. **Biological control:** propylene glycol acts as a natural inhibitor, preventing the formation of algae and biofilm that could obstruct the microchannels of the cold plates, whose fins are often less than 200 microns wide;
3. **Safety:** unlike ethylene glycol, propylene glycol is classified as low-toxicity, making it safe for use in indoor environments with personnel present;
4. **Performance:** although the addition of glycol slightly increases viscosity, thermal performance does not significantly differ from that of pure water.

Properties	Air	Water (technical grade)	PG25 (water/glycol)	Ratio (PG25/Air)
Density [kg/m³]	1.149	994,4	1017,5	~885x
Specific Heat [kJ/kg·K]	1.006	4.179	3.960	~3.93x
Thermal Conductivity [W/m·K]	0.026	0.621	0.493	~18,96x
Volumetric Heat Capacity [kJ/m³·K]	1.1156	4155,6	4029,3	~3611x

Note: properties calculated at 34°C

3.4 Fluids for two-phase direct-to-chip and immersion cooling

Unlike single-phase DTC, immersion cooling technologies and two-phase DTC require the exclusive use of dielectric fluids (synthetic oils or fluorinated hydrocarbons) to ensure the electrical insulation of components.

Although single-phase immersion cooling requires a thermodynamic compromise, as dielectric fluids exhibit specific heat capacity and thermal conductivity values approximately 50% lower than PG25, two-phase DTC and two-phase immersion cooling overcome this limitation by exploiting the latent heat of evaporation: these systems use fluids such as hydrofluoroolefins (HFOs) or fluorinated ketones, which evaporate at controlled temperatures (in the 30–60°C range).

However, fluid selection is now strictly constrained by sustainability: the industry is phasing out perfluorocarbons (PFCs) and hydrofluorocarbons (HFCs) due to their prohibitive GWP (Global Warming Potential), often exceeding 1,000, as well as increasing restrictions on PFAS (“forever chemicals”). The transition to next-generation fluids with GWP < 1 has therefore become an essential requirement for ESG compliance, ensuring that high computing density does not translate into regulatory risk or unsustainable atmospheric impact.

4 COOLING TECHNOLOGIES

4.1 Single-phase direct-to-chip (DTC)

Single-phase direct-to-chip (DTC) represents the current industry standard for AI and HPC.

Heat is removed through a copper cold plate with internal microchannels, mounted directly on the processor (CPU/GPU) using a thermal interface material. This technology involves the use of PG25, which remains in the liquid state at all times.

PLUS:

- **Targeted efficiency:** removes approximately 85% of the node's heat by acting directly on the most power-intensive components;
- **Versatility:** enables "hybrid" configurations, where liquid cools the processors and a reduced airflow manages secondary components (RAM, VRM, storage);
- **Retrofit:** it is the simplest solution to integrate into existing data centers with standard 19" rack layouts.

MINUS:

- **Leak management:** requires the use of high-quality drip-less quick disconnects (QD) and leak detection sensors within the chassis;
- **Cabling complexity:** the addition of manifolds and piping within the server reduces the space available for physical component maintenance.

4.2 Two-phase direct-to-chip (DTC)


In two-phase DTC, the fluid inside the cold plate undergoes a phase change, transitioning from liquid to vapor. It exploits the latent heat of evaporation to absorb energy without increasing the fluid temperature.

PLUS:

- **High thermal capacity:** ideal for next-generation chips exceeding 1000–1500 W, where single-phase flow would require excessive flow rates;
- **Thermal uniformity:** ensures that the entire chip remains at the same temperature (the liquid saturation temperature), eliminating dangerous thermal gradients.

MINUS:

- **Pressure control:** requires fully hermetic systems and extremely precise pressure control valves to manage the vapor;
- **Costs:** the use of specialized refrigerant fluids and the design of piping systems to manage vapor return significantly increase CAPEX
- **Fluid:** many of the options currently available on the market exhibit a high GWP.



4.3 Single-phase Immersion Cooling (IC)


The entire server is immersed in a tank containing a dielectric fluid, which is forcibly circulated toward an external heat exchanger. In this Immersion Cooling (IC) configuration, the fluid always remains in the liquid state.

PLUS:

- **100% liquid cooling:** Two-phase IC dissipates 100% of the server heat, including power supplies and networking components, eliminating the need for air cooling and consequently removing fans and noise.
- **Hardware protection:** The dielectric fluid protects the electronics from oxidation, humidity, and dust, reducing PCB failure rates.

MINUS:

- **Maintenance:** Removing a server requires complex draining and cleaning procedures.
- **Weight/layout:** Requires floors with very high load-bearing capacity and a complete redesign of the room layout, with horizontal tanks instead of vertical racks.



4.4 Two-phase immersion cooling (IC)

In this configuration, the dielectric fluid boils upon contact with the components. The vapor rises to the top of the tank, where it condenses on a water-cooled coil and returns by gravity.

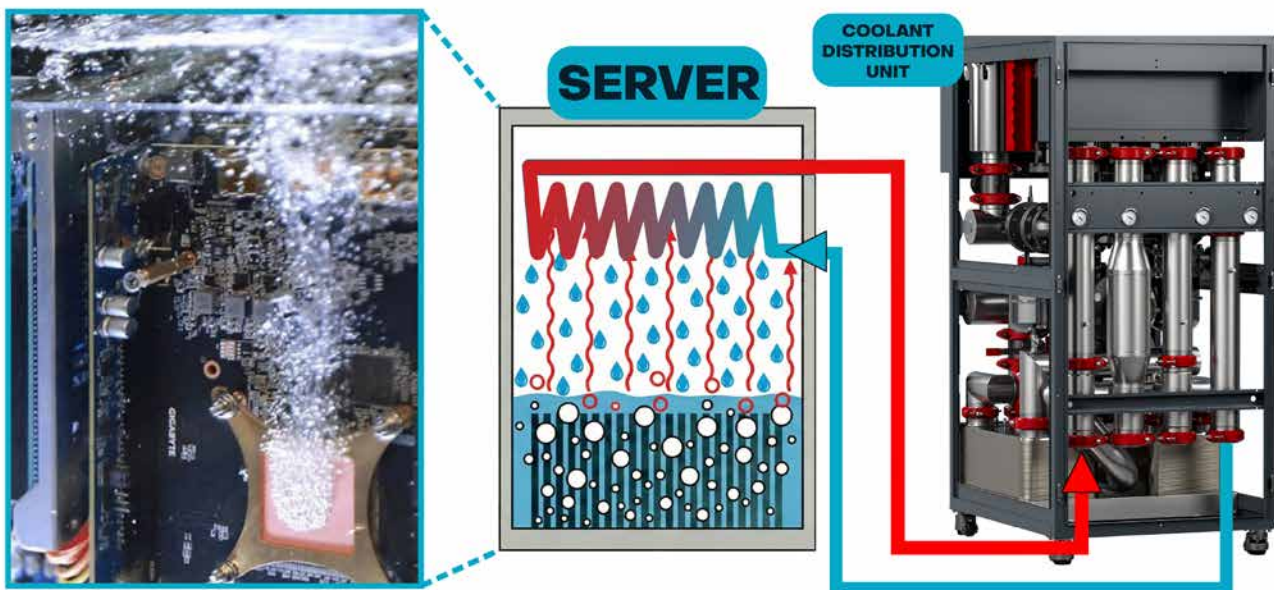
PLUS:

- **Efficiency:** enables the achievement of very low PUE values (down to 1.02), operating with primary circuit cooling water at very high temperatures;
- **No pumps in the secondary circuit:** fluid movement is driven by thermodynamics.

MINUS:

- **Fluid cost:** dielectric phase-change fluids are extremely expensive and subject to stringent environmental regulations;
- **Maintenance:** extracting a server requires complex draining and cleaning procedures;
- **Weight/layout:** requires floors with very high load-bearing capacity and a complete redesign of the room layout, with horizontal tanks instead of vertical racks;
- **Fluid:** many of the options currently available on the market exhibit a high GWP.

Technology	Supported density	Maintenance complexity	Integration	Application target
Single-phase DTC	40 - 120 kW	Medium	Good (hybrid)	AI training / HPC standard
Two-phase DTC	100 - 200 kW	High	Limited	Chip > 1500 W
Single-phase / two-phase immersion	100 - 300+ kW	High	Complex	Supercomputing



5 LIQUID COOLING ECOSYSTEM

5.1 Hybrid room architectures

Despite the efficiency of cold plates, direct-to-chip (DTC) cooling acts selectively only on high thermal density components (CPU, GPU, VRM, RAM). However, within a server there are numerous other components – inductors, capacitors, motherboard chipsets, storage units, and power supply units (PSUs) – that continue to dissipate heat into the internal chassis air. Typically, liquid cooling in a DTC system is capable of removing between 80% and 85% of the total rack thermal load. The remaining portion is dissipated into the air: therefore, the 2026 data center is not “liquid-only,” but hybrid.

It is important to emphasize that in rooms equipped with immersion systems, where servers are cooled 100% by liquid, adequate air control and room conditioning must still be ensured. This is necessary to dissipate the thermal load generated by operators, lighting, etc., as well as to maintain thermo-hygrometric conditions within a comfort range for personnel.

5.1.1 CDU + rear door heat exchangers (RDHx)

The RDHx is essentially a liquid-based radiator mounted on the rear door of the rack. It is not a DTC technology in the strict sense, but it acts as a "thermal filter" that neutralizes the heat generated by the servers before it exits the rack.

Questa è considerata l'architettura ibrida "pura". La CDU gestisce il circuito secondario e la dissipazione del calore sulle cold plates, mentre le rear-doors catturano il calore residuo in uscita dai server.

- **Advantage:** heat never enters the room. The rack is thermally neutral;
- **Ideal for:** rooms without hot/cold aisle containment or with low ceilings where effective airflow management is difficult.


The rear door represents a versatile technology that can also be deployed as a stand-alone solution for rooms with 30–40 kW per rack.

PLUS:

- **Retrofit:** does not require modifications to the nodes or installation of piping on the processors;
- **Room Neutral:** rende il data center indipendente dal sistema di condizionamento dell'aria ambiente (CRAC), poiché il rack non emette aria calda nella stanza.

MINUS:


- **Fan dependency:** servers must still maintain powerful internal fans to push air through the door, resulting in high energy consumption and noise levels;
- **Physical limit of air-based heat exchange:** above 40 kW, air cooling alone is no longer a viable option and must be integrated with liquid cooling.



5.1.2 CDU + Perimeter Units (CRAC/CRAH)

In this scenario, the CDU manages the secondary circuit and the cooling of the cold plates, while the perimeter units (CRAH) handle the residual load and ambient humidity control.


- **Advantage:** enables the reuse of existing infrastructure in retrofit scenarios;
- **Ideal for:** existing data centers adding AI clusters in dedicated areas of the white room.



5.1.3 CDU + Fanwall

In this scenario, the CDU manages the secondary circuit and the cooling of the cold plates, while the portion of the load dissipated into the air is handled by the fanwalls.

- **Advantage:** high energy efficiency, reduced footprint in the white space;
- **Ideal for:** large hyperscale DTC data centers designed ex novo.



5.2 The complete ecosystem of liquid cooling

The complete ecosystem of liquid cooling can be viewed as a three-level pyramid of heat exchanges and thermal flows:

1. **Server level:** the cold plates capture the heat generated by the chip and transfer it to the PG25;
2. **CDU level:** the plate heat exchanger within the CDU transfers heat from the secondary circuit (PG25) to the primary circuit (typically water / water-glycol);
3. **External level:** the chiller and/or the dry cooler dissipate the heat to the outdoor environment.

The proper integration of CDUs, air-based units, and external heat rejection systems redefines the energy profile of the data center, enabling cooling pPUE values below 1.1. The competitive advantage of this architecture lies in seasonal performance analysis: by operating with elevated primary circuit temperatures, it is possible to leverage free cooling for the vast majority of days throughout the year in many

climatic regions, limiting chiller operation to rare summer thermal peaks.

This approach enables the maximization of energy savings during cold and intermediate seasons, turning seasonal efficiency into a key driver for reducing operating expenses (OpEx) and achieving environmental sustainability targets.

5.3 Liquid distribution within the room

While in air-cooled data centers distribution occurs through underfloor pressurization, in liquid cooling distribution consists of a physical piping network that must be designed to ensure operational continuity and system balancing, as well as to minimize the risk of leaks. The key elements of the distribution circuit are as follows:

- **Distribution manifold:** manifolds that distribute the secondary circuit fluid among the various servers within the same rack. Manufactured in AISI 304/316L stainless steel with orbital/TIG welds to eliminate internal burrs that could detach and obstruct the cold plates;

- **Quick disconnects (QD):** connectors that provide the interface between the various critical components of the circuit. They represent the most stressed component throughout the data center lifecycle. The use of “non-drip” connectors with flat-face technology is essential to prevent liquid droplets from falling onto the servers.

6 COOLANT DISTRIBUTION UNIT



The Coolant Distribution Unit (CDU) represents the cornerstone of scalability for modern data centers. The current engineering challenge lies in maximizing heat exchange density while minimizing footprint.

Latest-generation units are capable of handling massive thermal loads, reaching capacities of up to 1400 kW within a footprint of just 900 mm × 1200 mm. Such power density is made possible by advanced engineering that ensures a 4 K thermal approach, enabling operation in full free-cooling mode for the majority of the operating year and drastically optimizing overall PUE.

The CDU acts as a separation barrier and management system between the two hydraulic domains that make up the data center:

1. **Primary loop (Facility Water System - FWS):** the building-side circuit, connected to external heat rejection systems, which typically operates with a 10°C deltaT;
2. **Secondary loop (Technology Cooling System - TCS):** the IT-side circuit, where the more “critical” fluid, PG25, circulates. Here, the CDU ensures pressure stability, fluid purity, and precise temperature control to prevent condensation.

6.1 CDU components



HEAT EXCHANGERS

Heat transfer between the primary and secondary fluids occurs through brazed plate heat exchangers made of AISI 316L stainless steel, characterized by a high heat transfer coefficient and excellent corrosion resistance. The heat exchangers are sized to provide a nominal 4 K approach temperature.

PUMP GROUP

The circulation of the secondary circuit fluid is handled by high-efficiency centrifugal pumps in an N+1 redundant configuration. The pumps are controlled via inverters and are therefore capable of modulating flow rate according to the instantaneous thermal load of the racks.

FILTER

The microchannels within the cold plates have dimensions on the order of 200 microns: it is therefore essential to provide appropriate filtration to prevent any possible fouling or obstruction. The CDU consequently integrates a redundant N+1 duplex filtration system:

- Filter cartridges are selectable down to 25 microns, ensuring the capture of even the smallest particles that could cause erosion or obstruction within the servers;
- The duplex configuration allows the cartridge to be isolated and replaced “hot,” without interrupting the operational continuity of the secondary circuit, thereby ensuring uptime even during routine maintenance activities.

PICV VALVE

Within the CDU, the pressure independent control valve (PICV) performs the dual function of flow regulator and pressure stabilizer, acting as the “precision regulator” of the primary loop:

- **Pressure independence:** unlike traditional control valves, the PICV maintains a constant flow rate to the heat exchanger, regardless of pressure fluctuations within the data center water network;
- **DeltaT optimization:** by dynamically regulating water flow according to the actual thermal load of the servers, the PICV ensures that the water leaving the primary circuit is at the highest possible temperature. This is essential for the efficiency of district heating systems or chillers operating in free-cooling mode.

CONFIGURABLE CONNECTIONS

To accommodate both new buildings (greenfield) and the refurbishment of existing rooms (retrofitting), the CDU is designed with configurable top and bottom connections. This configuration allows piping from either the ceiling or the raised floor to be connected without requiring structural modifications to the unit or the addition of joints or fittings during installation.

CHARGE CONNECTION

To facilitate filling of the secondary loop with PG25, the CDU features a charging connection on the secondary circuit designed for connection to a filling pump. In some configurations, the pump can be directly integrated within the CDU frame.

6.2 Regulation logic

The complexity of liquid cooling lies in the extreme dynamic nature of the loads. Compared to traditional enterprise workloads, which exhibit predictable and gradual variations, the AI ecosystem introduces unprecedented thermal volatility.

LOAD DYNAMICS DURING TRAINING

During the training phases of a large language model (LLM), servers do not operate at a constant regime. The thermal load follows the progression of computational “batches”:

- **Instantaneous peaks:** when the cluster begins processing a dataset, rack power consumption can rise from 10 kW (idle/waiting) to 80–100 kW within a few seconds;
- **Cyclic oscillations:** during training, continuous cycles of “compute” (100% load) are observed, followed by very brief “sync/communication” pauses (load drop), creating a “sawtooth” thermal profile.

Unlike air-based systems, where the air mass within the room acts as a partial “thermal flywheel,” in liquid cooling heat is transferred instantaneously to the fluid. Without a CDU capable of modulating flow rate, these oscillations would cause thermal stress on the processor dies, shortening their service life or triggering thermal throttling phenomena that slow down the entire training process.

PID CONTROL AND FLOW MODULATION

To manage this variability, the core of the control strategy is based on PID algorithms acting on the pump inverters (VFDs).

- **Differential pressure management:** the CDU continuously monitors the pressure difference between supply and return in the secondary circuit, with the objective of maintaining a constant deltaP at the racks, enabling the system to respond instantaneously to the opening or closing of local valves on server nodes during computational peaks;
- **Thermal modulation:** by adjusting pump speed and the PICV valve on the primary circuit, the CDU stabilizes the supply temperature at the target setpoint, compensating for rapid return temperature fluctuations caused by AI training cycles.

DEW POINT CONTROL AND MANAGEMENT

One of the greatest risks in a hybrid data center room is condensation formation: the CDU therefore integrates ambient humidity and temperature sensors, calculating the room dew point in real time and ensuring that the supply fluid temperature remains at least 3–5°C above the dew point.

6.3 NVIDIA server compliance

For a designer, compliance with chip vendor standards is a prerequisite to ensuring the reliability of the IT investment. NVIDIA Blackwell (GB200) architectures impose stringent parameters on CDUs:

- **Filtration:** NVIDIA requires stringent filtration to protect GPU microchannels. The configuration of our units with 25-micron cartridge filters exceeds these standards, preventing erosion of internal components;
- **Flow stability:** the CDU must handle sudden variations without pressure oscillations exceeding 5%, a critical specification for HGX/GB200 systems;
- **Communication protocols:** native support for Modbus TCP/IP, BACnet, or SNMP, enabling full telemetry integrated into the cluster management system;
- **Redundancy:** N+1 redundant pumping system and redundant controllers to meet High Availability requirements.

7 EFFICIENCY, SUSTAINABILITY, AND HEAT RECOVERY

7.1 Beyond PUE: the definition of ERF

Power Usage Effectiveness (PUE) has for years been the sole benchmark for data center efficiency. However, with the advent of liquid-cooled data centers, the industry is now facing a dual trend: PUE values are increasingly approaching unity, while both operating temperatures and power densities within data centers continue to rise, making heat recovery both straightforward and necessary.

At this stage, PUE alone is no longer sufficient to describe the “true” impact of a data center and must be complemented by a KPI capable of quantifying its “circular” value, namely how much of the generated heat is recovered and reused: this leads to the definition of the Energy Reuse Factor (ERF).

ERF = Reused Energy / Total Data Center Energy

While PUE merely measures how much energy is “wasted,” ERF quantifies how much energy is reinjected into the system: the paradigm shift is clear, moving from a “heat-as-a-waste” model to a “heat-as-a-source” approach. Therefore, the greater the share of energy that is recovered and reused, the lower the ERF value will be.

7.2 Heat recovery strategies

In liquid cooling, recovering the heat generated by the IT load is simpler than in air-based systems. In fact, the return temperature of the secondary fluid reaches values close to or even exceeding 40°C, making both direct reuse and integration into a very high-temperature heat pump feasible.

DISTRICT HEATING

The primary circuit of the CDU is connected to a very high-temperature heat pump, which supplies the district heating network.

A 10 MW AI cluster can heat approximately 2,000–3,000 modern homes: in many European countries, this enables the data center to obtain government incentives or significantly reduce CO₂ emission taxes.

INDUSTRIAL PROCESSES / AGRICULTURE 4.0

Beyond district heating, constant medium-temperature heat available 24/7 lends itself to a wide range of additional applications:

- **Hydroponic greenhouses:** maintenance of optimal temperatures for intensive crops, reducing heating costs;
- **Drying or preheating systems in industrial processes** requiring hot water for washing or chemical preparation stages;
- **Swimming pool heating.**

OPTIMIZATION OF THE LOCAL INFRASTRUCTURE

In the absence of external networks, the heat can be used for:

- Improve the efficiency of the data center itself;
- Heat administrative offices / warehouses;
- Preheat domestic hot water (DHW) for adjacent buildings.

7.3 Regulatory framework

The adoption of liquid cooling and the implementation of strategies to maximize the **ERF (Energy Reuse Factor)** are now fundamental pillars for compliance with new international directives and for access to capital markets.

THE EUROPEAN EED DIRECTIVE

The revision of Directive EED 2023/1791 (Energy Efficiency Directive) has introduced unprecedented reporting and transparency obligations for the sector:

- **Heat recovery obligation:** Delegated Regulation (EU) 2024/1364 requires data centers with a nominal capacity above 1 MW to conduct cost-benefit analyses for the utilization of waste heat through district heating networks or other industrial applications, with a minimum ERF of 0.20;
- **European database:** from 2024, operators of data centers above 500 kW are required to publish their KPIs annually, including PUE, WUE, and ERF, within the European platform for data center efficiency monitoring;
- **EU green taxonomy:** liquid cooling represents one of the “enabling technologies” that facilitate compliance with emission thresholds required to classify an investment as “sustainable” under the technical screening criteria of the EU Taxonomy (Regulation 2020/852).

DECARBONIZATION OF SCOPE 2 AND SCOPE 3

The integration of direct-to-chip systems directly impacts the corporate decarbonization strategy defined by the **GHG Protocol (Green House Gas Protocol)**:

- **Scope 2 reduction (indirect energy emissions):** the intrinsic efficiency of liquid drastically reduces the energy required for cooling (PUE below 1.1), lowering emissions associated with purchased electricity;
- **Scope 3 reduction (value chain and circularity):** within a circular economy framework, the integration of heat recovery systems enables the data center to evolve from an energy-intensive user into an active supplier of thermal energy to the community, generating so-called “avoided emissions.” Although, according to the GHG Protocol, these do not mathematically subtract tonnes of CO₂ from the company’s Scope 3, their disclosure under the CSRD (Corporate Sustainability Reporting Directive)

demonstrates a regenerative business model and industrial symbiosis.

This circular approach reduces regulatory transition risks, significantly improves ESG ratings, and facilitates access to sustainable finance instruments such as green bonds, positioning the infrastructure as a priority asset for institutional investors focused on long-term sustainability.

8 OPERATIONAL CHALLENGES, MAINTENANCE, AND SAFETY

8.1 Chemical management and monitoring of PG25

The chemical stability of PG25 is essential to prevent galvanic corrosion and biofouling: it is therefore necessary to perform at least semi-annual sampling of the secondary circuit fluid in order to monitor:

- **Glycol level:** to ensure that freeze protection and thermal properties remain consistent;
- **Alkalinity reserve (pH):** a decrease in pH indicates glycol oxidation, which can become acidic and attack metals. The ideal pH value should range between 8.0 and 9.5;
- **Inhibitor concentration:** copper and stainless steel passivating agents must remain active.

Proper management of secondary circuit refilling is essential: any fluid top-up must be carried out using a pre-diluted industrial-grade mixture, while the addition of untreated tap water would compromise the electrical conductivity and purity of the loop.

8.2 Routine maintenance of the CDU

The presence of microchannels in the cold plates (with gaps below 200 microns) makes filtration the most critical maintenance operation.

- **Duplex basket filters in the CDU:** maintenance of the 25-micron filters is performed by monitoring differential pressure. An increase in deltaP indicates cartridge clogging. Thanks to the duplex configuration, hot maintenance can be carried out by isolating one filter at a time through the shut-off valves, without interrupting machine operation;
- **Degassing:** air trapped within the circuit can cause localized erosion; it is therefore essential to regularly verify the proper operation of automatic air vent valves.

8.3 Server maintenance

Maintenance of a modern direct-to-chip (DTC) system is designed to be easy and fast, while ensuring operational continuity.

HOT-SWAP OPERATIONS

The core of maintainability lies in the UQD (Universal Quick Disconnect) fittings. These connectors are specifically engineered to be intrinsically “non-drip,” minimizing the risk of PG25 leakage onto servers or, more generally, within the white space, while simultaneously preventing air entrapment

within the circuit. Furthermore, thanks to rack modularity and manifolds equipped with shut-off valves for each slot, targeted intervention is possible by isolating only the server undergoing maintenance, without affecting flow rate or pressure across the remaining nodes.

LEAK PREVENTION

Preventive maintenance includes inspection of internal chassis hydraulic paths and seals in order to prevent any potential fluid leaks:

- **Leak detection:** liquid-cooled racks integrate cable or point sensors positioned at the bottom of the chassis, capable of sending an alarm signal and shutting down the server in the event of abnormal humidity values;
- **Tube bending:** during maintenance, it is essential to verify that internal flexible hoses (FEP/EPDM) have not undergone twisting or kinking that could increase local hydraulic resistance;
- **Quick disconnect seals (O-rings):** O-rings must be periodically inspected. The use of high-quality seals reduces maintenance requirements, but frequent insertion/removal cycles may require light lubrication with PG25-compatible products to prevent rubber tearing that could lead to micro-leaks.

9 FUTURE FRONTIERS – 800 V DATA CENTERS

9.1 Current state: 400 V AC distribution

Today, in data centers, distribution typically occurs at 400 V AC within the room and at 48 V DC at the rack level. However, the advent of AI clusters with densities exceeding 150 kW per rack places significant electromechanical stress on the current architecture based on 48 V copper busbars (power bars).

Considering a 100 kW rack, 48 V busbars require a current exceeding 2,000 A. To handle such current levels, either multiple parallel busbars or a single oversized busbar must be used: in both cases, this results in high costs, significant physical dimensions, and increased cabling complexity. Furthermore, it must be considered that Joule losses along the internal rack distribution lines, which generate an additional “parasitic” thermal load, increase quadratically with rising current: increasing the voltage therefore allows these losses to be minimized.

9.2 800 V DC distribution

Borrowed from high-performance electric vehicle (EV) engineering, 800 V direct current distribution is emerging as the definitive solution for next-generation data centers. The advantages are multiple:

- **Busbar cross-section reduction:** with a voltage of 800 V instead of 48 V, the current required to supply a 100 kW rack drops to approximately 120 A. This enables the use of lighter and more flexible conductors, optimizing internal chassis space and facilitating maintenance operations;
- **Transmission efficiency:** since line losses decrease with the square of the voltage, transitioning from 400 V to 800 V increases electrical distribution efficiency beyond 99%, drastically reducing the heat generated by the cables;
- **Simplification of the conversion chain:** the HVDC (High Voltage Direct Current) architecture eliminates intermediate AC/DC rectification and transformation stages. Fewer conversion stages mean fewer components subject to failure and a lower total amount of heat to dissipate.

9.3 Liquid-cooled PSU

The extreme power density of new 800 V power supply units (PSUs) renders forced-air ventilation obsolete. As a result, next-generation PSUs integrate cold plates for cooling, directly connected to the secondary circuit managed by the CDU: liquid cooling is no longer limited to CPUs and GPUs, but extends to nearly all rack components. This shift increases the percentage of load removed by liquid

to approximately 95–98%, thereby drastically reducing the need for air cooling.

9.4 Challenges of 800 V: safety and environmental management

Operation at 800 V DC introduces stringent safety requirements:

- **Arc fault prevention:** at such voltage levels, faults can trigger sustained electric arcs. The use of connectors with safety interlocks and ultra-fast protection systems capable of isolating the fault within milliseconds is essential;
- **Condensation prevention:** the proximity between 800 V conductors and liquid circuits requires impeccable dew point control. The CDU acts as a safety guarantor: by modulating the supply temperature to prevent condensation, it mitigates the risk of catastrophic short circuits in an HVDC environment.

10 CONCLUSION

In 2026, it can be stated with certainty that liquid cooling has definitively moved beyond the HPC niche to become the only technical and infrastructural paradigm capable of sustaining the AI revolution.

With processors having long surpassed the “critical” 1000 W threshold, air is no longer a viable path for data center cooling: the transition to liquid cooling is no longer a simple option, but an imperative.

This revolution also marks the integration of sustainability into data center KPIs: heat recovery becomes both a regulatory requirement and a financial lever, shifting the paradigm from heat-as-waste to heat-as-source.

In this way, the data center evolves from being an “isolated bunker” into a true “urban energy ecosystem,” silent and integrated, capable of heating entire communities.

11 APPENDIX A

11.1 Chemical management of the fluid and material compatibility

FILLING WATER SPECIFICATIONS

Prior to mixing with propylene glycol, the water used must meet extremely high purity requirements. The use of tap water, even if softened, is strictly prohibited due to the presence of chlorides and sulfates.

Typical properties of PG 25 and PG 55 heat transfer fluids:

Characteristic	Performance
Appearance	Clear and particulate free
Fluid pH	8.0 – 10.5 The pH of the fluid is dependent on corrosion inhibitor formulation and may be lower when using organic acid technology (OAT)
Unadjusted reserve alkalinity	4 mL, based on the diluted fluid Measure of the buffering capacity of the fluid
Copper-copper hex	<2 ppm
Iron	<2 ppm
Total hardness	<20 ppm High hardness values indicate the use of poor quality water
Chlorides	<5 ppm High chloride values indicate the use of poor quality water
Sulphates	<10 ppm High sulfate values indicate the use of poor quality water

PG25 CHARACTERISTICS

The operating fluid must be a mixture of industrial-grade propylene glycol with a pre-formulated multi-metal corrosion inhibitor package.

Functions of the inhibitors:

- Copper passivators (azoles):** create a monomolecular film on copper cold plate surfaces to prevent oxidation and ionic dissolution;
- pH buffers:** maintain the fluid in an alkaline environment (pH between 8.5 and 10.5) to neutralize acidic by-products resulting from the natural degradation of glycol over time;
- Steel/stainless steel inhibitors:** protect manifolds and CDU components from pitting corrosion.

Risks of degradation (Glycolysis)

When exposed to elevated temperatures and oxygen, propylene glycol can degrade into organic acids (glycolic acid, formic acid). This lowers the pH and transforms the fluid into a

corrosive agent. Monitoring the alkalinity reserve is the only way to predict this phenomenon.

MATERIAL COMPATIBILITY

To ensure a 20-year service life, the secondary loop must comply with stringent requirements regarding galvanic compatibility: interaction between metals with distant electrochemical potentials is the primary cause of failures in liquid-to-chip systems.

Approved materials:

- **Metals and metal alloys:** Copper, Brass with <15% zinc, Stainless steel, Nickel, high nickel alloys, Chromium, Titanium, B-Ni-6, BCuP-2, BCuP-3, BCuP-4, BCuP-5, TF-H60F;
- **Elastomers, plastics and other materials:** EPDM, Viton A, Viton GF, Viton ETP, FEP, PTFE, PP, HDPE, PEEK, Loctite 567.

Prohibited materials:

- **Aluminum (including anodized):** extremely sensitive to galvanic corrosion in the presence of copper;
- **Galvanized steel:** zinc reacts with glycol inhibitor packages, forming sludge and solid deposits;
- **Tin/lead solder joints:** may be chemically attacked by glycol inhibitors;
- **Sulfur-containing EPDM:** may release compounds that attack copper.

12 APPENDIX B

12.1 Operational commissioning and testing protocol

This protocol defines the standard procedures for Testing and Commissioning (T&C) of Coolant Distribution Units and the associated hydraulic infrastructure. The objectives are to validate mechanical integrity and the dynamic response of control systems prior to IT load integration, as well as to ensure overall circuit cleanliness.

PHASE 1: PNEUMATIC CIRCUIT TESTING

Before filling the secondary circuit with PG25, the system must be pressure-tested to identify any potential leaks.

- **Procedure:** isolate the CDU from the racks using the shut-off valves; – pressurize the secondary circuit with dry nitrogen to a pressure equal to 1.5 times the design operating pressure (typically 6–8 bar);
- **Duration: 24 hours;**
- **Acceptance criterion:** pressure drop <1% after ambient temperature compensation.

Note: the use of nitrogen is mandatory instead of compressed air in order to prevent the introduction of moisture and oily contaminants into the loop.

PHASE 2: FLUSHING AND PASSIVATION

Circuit flushing is a critical operation: if not properly removed, welding residues, metallic particles, and machining oils can destroy the cold plates in a very short time. Closed-loop flushing protocol:

1. **Server bypass:** ensure that all racks/servers are isolated from the circuit. The flushing fluid must never pass through the cold plates during this phase;
2. **Detergent solution:** use demineralized water supplemented with a dispersing detergent specifically formulated for multi-metal circuits;
3. **Flushing velocity:** the CDU pumps must operate at maximum frequency to ensure a fluid velocity of at least 1.5 m/s, required to carry heavy debris toward the CDU filters;
4. **Filter inspection:** monitor the deltaP across the basket filters. Clean the cartridges approximately every 2 hours until no further debris accumulation is observed;
5. **Passivation:** after flushing, a passivating agent must be applied to create a protective film on internal copper and stainless steel surfaces, preventing immediate oxidation prior to final filling with PG25.

PHASE 3: FINAL FILLING AND DEAERATION

Filling with the operating fluid (PG25) must be carried out in a manner that prevents pump cavitation.

- **Filling procedure:** use the auxiliary pumping unit (make-up pump) to charge the system from the lowest point. The CDU must be equipped with a charging connection for the make-up pump or integrate the pump directly within the unit.
- **Degassing cycle:** operate the CDU pumps at variable speed (20%–80%–20% cycles) to promote the detachment of air bubbles from the piping walls.
- **Vent monitoring:** manually verify each air vent valve at the high points of the rack manifold. A properly deaerated system should be acoustically “silent.”

PHASE 4: VALIDATION OF CONTROL LOGICS

PID cycle test

- **Procedure:** simulate IT load variation by opening and closing the shut-off valves of multiple racks.
- **Verification:** the CDU must modulate pump speed to maintain the programmed deltaP setpoint with a maximum deviation of 5%.

Dew point control test

- **Procedure:** use a portable steam generator near the CDU hygrometric sensors to artificially increase the detected humidity.
- **Expected response:** the CDU must force the opening of the primary-side mixing valve to raise the supply temperature above the calculated dew point, while simultaneously activating the visual/remote alarm.

Pump failover

- **Procedure:** simulate an electrical fault on Pump 1 (magnetothermal trip).
- **Verification:** Pump 2 (in standby) must start within 5 seconds, ensuring flow continuity without the supply temperature exceeding the critical alarm threshold.

PHASE 5: SITE ACCEPTANCE TEST (SAT) AND HAND-OVER

The final phase involves the signing of the commissioning documentation and the handover of telemetry data.

- **BMS integration:** verify that all Modbus/BACnet points are readable by the central data center management system.
- **Verification of nominal flow rates:** measurement with an external ultrasonic flow meter on each rack branch to confirm that hydraulic distribution is balanced.

- **Final documentation:** delivery of the hydraulic “as-built” documentation, the initial fluid chemical analysis logs, and the sensor calibration certificates.



Viale Spagna, 31/33 - 35020 Tribano (Padova) - Italy

Ph. +39 049 9588511 - Fax +39 049 9588522

VAT # IT02191431200

info@hiref.it - www.hiref.it

